

ChatGPT Took My Tech Entrepreneurship Exam. Here's What Happened.

Shyam Kumar (Rensselaer Polytechnic Institute)

KEYWORDS: Teaching Methods, How to start a business.

The test measures students' knowledge of both established principles and managerial "soft" skills. Check out the results -- and test your answers against ChatGPT on seven questions.

We all want to understand how Large Language Models (LLMs), most prominently ChatGPT, can help people in different fields make better decisions. In theory, ChatGPT could serve as a valuable tool and partner in decision-making if it has sufficient domain-specific knowledge. Studies have examined the [performance of ChatGPT](https://theamericangenius.com/housing/real-estate-tech/chatgpt-passing-exams/) on various standardized exams in professions with relatively well-identified and certifiable bodies of knowledge, and the results have been nothing short of astounding. ChatGPT has performed extremely well, demonstrating a high level of competency in these different fields.

These developments bring up interesting questions: How much entrepreneurship does Chat GPT know? And how comprehensive is its knowledge in a field that is less technical in nature and not defined by a discrete body of knowledge? These questions are very salient to me in my work as a professor at Rensselaer Polytechnic Institute's Lally School of Management. So to test these questions, I decided to put ChatGPT through the final exam of my Principles of Technological Entrepreneurship class as an experiment to gauge its performance in a domain that is not only broad in terms of knowledge boundaries, but also inherently driven by soft skills.

This exercise is not just of interest to professors like me, or to our students. Rather, it has significance for a variety of entrepreneurship practitioners for a couple of reasons. First, as we increasingly turn toward using

LLMs such as ChatGPT as "co-pilots" in our lives, we need to know whether they are reliable sources of information, and whether they can be trusted, at least to some degree, as we engage with them. Second, many sources in the popular press have touted the [benefits of LLMs](https://www.forbes.com/sites/aytekintank/2024/05/29/5-chatgpt-prompts-to-find-profitable-startup-ideas/)

(<https://www.forbes.com/sites/aytekintank/2024/05/29/5-chatgpt-prompts-to-find-profitable-startup-ideas/>) in helping budding entrepreneurs generate and evaluate ideas, and ChatGPT also has a module that works with Y Combinator to provide startup advice. This makes it all the more critical to know if ChatGPT can be relied upon to guide entrepreneurs not just in idea generation, but also at different stages of their entrepreneurial journey.

About My Class, and the Test

The Technological Entrepreneurship class that I teach can best be described as a class designed to give students a well-rounded exposure to entrepreneurship principles. It covers topics ranging from how entrepreneurship is different from managing established business ventures; to lean startup principles; to related concerns such as legal issues; and the relevance of ["Crossing the Chasm"](https://www.amazon.com/Crossing-Chasm-3rd-Disruptive-Mainstream/dp/0062292986) (a widely-used model of technology adoption), and the role of business models. The course delves into finance and marketing for startups, and culminates with the challenges of scaling and exit. The course also has a life cycle approach built into it (i.e. from a startup idea, to various stages of evolution) and can be considered a capstone entrepreneurship class.

The exam I administered to ChatGPT was intended to test students' critical thinking and understanding of the principles covered in class. The tone is theoretical and conceptual, and arguably an experienced entrepreneur would do well on the exam, even without taking the



course and going through the readings and class slides, given the emphasis on higher-order skills. The course did not adopt a textbook; it was based on a collection of slides and readings that I put together. This is important, because it implies ChatGPT did not have access to any previous test banks that it could have potentially been trained on.

Methodology

My experiment used Chat GPT 3.5's May 24, 2023 version. I began with a fairly simple prompt to ChatGPT: "I would like you to answer the below multiple-choice questions." I then cut and pasted each multiple-choice question and documented the responses.

The exercise was designed to mimic a test taking experience -- i.e. to gauge how a test taker would respond when given one chance to address each question. Hence rather than prompt ChatGPT further when responses were incorrect with a Chain-of-Thought approach, or probe whether it modified responses when given a question multiple times and test its AI capabilities, I simply noted down its first response. As such, if an entrepreneur starts with a good prompt and engages in further interactions with ChatGPT, this exercise reveals that ChatGPT is likely to provide useful guidance.

The Questions: Test Your Answers Against ChatGPT

Here are seven of the 50 questions on the test. If you'd like, try to answer these yourself. [Then click on this link](#)

(<https://eiexchange.com/api/submissions/1126/documents?file=How%20ChatGPT%20Answered.pdf&download=1>) to download the document that explains how ChatGPT answered, and its rationale for each answer. (Note: ChatGPT did not get a perfect score.)

Q1: Entrepreneurship as a concept

We defined entrepreneurship as "the pursuit of opportunity without regard to resources currently controlled." The definition contrasts entrepreneurship with managing established corporations in all the following respects, except (i.e. which of the below statements is false):

- Entrepreneurs are not constrained by existing resources, like managers of established firms
- Entrepreneurs are not constrained by social

- commitments to employees etc., like managers of established firms
- Entrepreneurs embrace greater risk-taking unlike managers of established firms
- Entrepreneurs operate in environments with high ambiguity compared to managers of established firms

Q2: The role of customer money and business plans

We learned that the best money for entrepreneurs comes from customers. Customer money, as opposed to external money, is preferable because:

- Customer money ensures the venture is focused on a real problem
- External money takes time to raise
- External money may dilute founders and lead to loss of control
- All of the above are true

Q3: The lean startup

Some limitations and potential risks of the lean startup approach include:

- It works only in startups and not in established corporations
- MVPs that learn at the expense of customers can harm reputation
- It cannot be implemented in rapidly changing industries
- It can only be implemented for technical products and not consumer products

Q4: Implications of business models

Changes in business model parameters can have cascading effects, and it is important to think through the implications when making changes. Which of the following statements would you disagree with and is most likely false:

- Business model changes can affect IT infrastructure costs
- Business model changes can affect logistics and value chain activities
- Business model changes can affect metrics that are monitored
- Business model changes can affect division of equity between founders

Q5: Pitching to different investors

When pitching to different investors, entrepreneurs need to modify their pitches to align with the incentives/risk factors of each investor. When compared to equity holders, pitches to debtors and banks should emphasize:

- Tangible assets, upside potential, home runs
- Intangible assets, upside potential, home runs
- Tangible assets, steady cash flows, limited down side risk
- Intangible assets, steady cash flows, limited down side risk

Q6: Choosing an appropriate VC

Not all VC money is equal. A VC manager is least preferable and attractive from the perspective of a startup when she/he:

- Has industry connections
- Is on the boards of several other startups
- Is a senior partner in the VC company
- Is experienced with different types of exits

Q7: Where entrepreneurs come from

Entrepreneurs come from several different sources and backgrounds. Data suggests the highest percentage of entrepreneurs come from:

- Other VC backed firms
- Universities, and are fresh out of college
- From large established firms
- Government

Chat GPT's Overall Score

ChatGPT scored a solid B+ on the exam based on its responses to 50 multiple choice questions. The explanations it provided suggest it has accumulated a robust repository of the core principles of entrepreneurship and has mapped the field well, validating its usefulness as a copilot and a potential information source. Equally importantly, the performance of ChatGPT on the exam reveals that soft skills are no longer the domain of humans alone, and that machines are attaining an impressive degree of mastery over the principles underlying these skills. Hence the advantage of humans over machines lies in generalizing these soft skills, and in judging how to apply them in different contexts taking into account complex interactions, including in the field of

entrepreneurship.

Some Caveats

The score on the test and the grade should be considered indicative, and not definitive in the sense of being a final grade awarded to ChatGPT, since it is constantly learning. In the future, conducting this exercise with later versions of ChatGPT would be worthwhile to document its improvements. Also worthwhile would be to pose each question multiple times to see the variation in answers at the same temperature setting, as well as different temperature settings. Finally, to the best of my knowledge, all questions on this test are original, since it was designed by me and did not use any off-the-shelf questions. Hence it is unlikely that ChatGPT was trained directly on any of the exam questions.

Also worth noting: this exercise used Chat GPT 3.5's May 24, 2023, version, with the default temperature setting. With later versions, in principle, ChatGPT's performance should only improve. But the main takeaway of this exercise should remain valid, which is that ChatGPT can be reasonably trusted as a tool and copilot for information related to entrepreneurship principles, including those that draw on higher-order soft skills.

Takeaways

In an earlier white paper, "Would Chat GPT get a Wharton MBA?", Christian Terwiesch noted some excellent implications that were confirmed in this experiment. In the interest of space, I will not repeat them, but here are some added insights:

ChatGPT 'understands' general entrepreneurship principles

Its performance on the test showed that ChatGPT has mapped the general principles underlying the field of entrepreneurship quite well and has captured and synthesized the ideas fairly comprehensively. Hence it is a good go-to resource and tool to gain domain knowledge and information about entrepreneurship with minimal search costs for entrepreneurs and those interested in pursuing business ownership. ChatGPT and LLMs can be useful in spurring conversations and chain-of-thought interactions on specific topics related to entrepreneurship, such as what makes for a good VC, or how to apply lean startup principles etc. However, scholars should continue to test and

document ChatGPT's depth of knowledge in sub-fields of entrepreneurship, such as entrepreneurial finance. This exercise is the first of its kind in that it attempts to assess the domain knowledge of ChatGPT, and we could gain from conducting similar exercises in other areas related to entrepreneurship.

Machines are beginning to grasp 'soft skills'

While Christian Terwiesch's exam mostly tested "hard concepts" and their application, my test had a greater emphasis on conceptual thinking related to soft skills. ChatGPT's performance on my exam suggests that it has good knowledge of soft skills in the field of entrepreneurship, and the underpinnings of those skills. We generally associate machines with mastery over hard skills such as mathematical calculations, but ChatGPT's performance on this exam shows a clear departure from that perspective.

Moreover, ChatGPT has several advantages in this realm. Unlike humans, it will remember these soft skill-related principles forever, and its repository of these principles is only likely to grow over time. We may therefore need to revisit the premise that soft skills differentiate humans from machines, and frame it more carefully.

Specifically, ChatGPT will be better able to access and retrieve the broader principles related to soft skills than a human who may have taken a Technological Entrepreneurship course, say 20 years ago.

Humans have the edge when judgment is needed

Compared to a human who is continually utilizing those skills, ChatGPT is likely to be behind and be unable to compete in its ability to generalize those skills to newer, emerging situations and apply them effectively. Hence, what is likely to differentiate humans from machines is not knowledge of principles underlying soft skills per se, but the ability to apply those skills when there is uncertainty and incomplete information, and where judgment is ultimately needed to assess the impact of complex interactions (e.g., Go ahead and experiment with an MVP despite some reputation costs). This conclusion echoes Agarwal, Gans and Goldfarb (2018), who base it on the predictive capabilities of machines, and it similarly applies to generative AI domains where machines are able to retrieve information and "reason" using principles related to soft skills, such as LLMs.

ChatGPT can make tests better for humans

For entrepreneurship scholars and educators in general, one takeaway is that it might be useful to administer a test intended for students to ChatGPT first, then study its reactions and explanations. This will give a sense of the extent of knowledge related to a particular domain that is out there in the public sphere, and that can be synthesized by an AI tool such as ChatGPT to answer a particular question. It can also be very useful to see the reasoning of an agent other than oneself when answering a question. This can reveal blind spots and help refine the question, leading to a better testing experience for students.

Conclusion

While there has been much discussion of the impact of LLMs, few investigations focused specifically on their depth of domain knowledge and reasoning, especially in management-related entrepreneurship soft skills. Its performance remains impressive, and it underscores why it should be adopted as a tool for advancing entrepreneurship education and practice.

References

Agarwal, A., Gans, J., Goldfarb, A. 2018. *Prediction Machines*. Harvard Business Review Press: Boston, MA.

Terwiesch, C. 2023. "Would Chat GPT Get a Wharton MBA? A Prediction Based on Its Performance in the Operations Management Course." Mack Institute for Innovation Management at the Wharton School, University of Pennsylvania.

Weij, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., Zhou, D. 2022. "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models". [arXiv:2201.11903v6](https://arxiv.org/abs/2201.11903v6) (https://arxiv.org/abs/2201.11903v6) [cs.CL]

Acknowledgement

I would like to express my thanks to my wife, Shimantika Kumar, for her input, edits, and invaluable feedback on this document. The usual disclaimer applies.